

METHOD AND APPARATUS FOR GENERATING A PHASE DEPENDENT CONTROL SIGNAL

TECHNICAL FIELD

This invention relates to generating a control signal and, more particularly, to generating a control signal based on the phase relationship between two input clock signals, and to memory devices and computer systems using such control signal generators.

BACKGROUND OF THE INVENTION

Conventional computer systems include a processor (not shown) coupled to a variety of memory devices, including read-only memories ("ROMs") which traditionally store instructions for the processor, and a system memory to which the processor may write data and from which the processor may read data. The processor may also communicate with an external cache memory, which is generally a static random access memory ("SRAM"). The processor also communicates with input devices, output devices, and data storage devices.

Processors generally operate at a relatively high speed. Processors such as the Pentium® and Pentium II® microprocessors are currently available that operate at clock speeds of at least 400 MHz. However, the remaining components of existing computer systems, with the exception of SRAM cache memory, are not capable of operating at the speed of the processor. For this reason, the system memory devices, as well as the input devices, output devices, and data storage devices, are not coupled directly to the processor bus. Instead, the system memory devices are generally coupled to the processor bus through a memory controller, bus bridge or similar device, and the input devices, output devices, and data storage devices are coupled to the processor bus through a bus bridge. The memory controller allows the system memory devices to operate at a clock frequency that is substantially lower than the clock frequency of the processor. Similarly, the bus bridge allows the input devices, output devices, and data storage devices to operate at a frequency that is substantially lower than the clock

frequency of the processor. Currently, for example, a processor having a 300 MHz clock frequency may be mounted on a mother board having a 66 MHz clock frequency for controlling the system memory devices and other components.

Access to system memory is a frequent operation for the processor. The
5 time required for the processor, operating, for example, at 300 MHz, to read data from or write data to a system memory device operating at, for example, 66 MHz, greatly slows the rate at which the processor is able to accomplish its operations. Thus, much effort has been devoted to increasing the operating speed of system memory devices.

System memory devices are generally dynamic random access memories
10 ("DRAMs"). Initially, DRAMs were asynchronous and thus did not operate at even the clock speed of the motherboard. In fact, access to asynchronous DRAMs often required that wait states be generated to halt the processor until the DRAM had completed a memory transfer. However, the operating speed of asynchronous DRAMs was
15 successfully increased through such innovations as burst and page mode DRAMs, which did not require that an address be provided to the DRAM for each memory access. More recently, synchronous dynamic random access memories ("SDRAMs") have been developed to allow the pipelined transfer of data at the clock speed of the motherboard. However, even SDRAMs are typically incapable of operating at the clock
20 speed of currently available processors. Thus, SDRAMs cannot be connected directly to the processor bus, but instead must interface with the processor bus through a memory controller, bus bridge, or similar device. The disparity between the operating speed of the processor and the operating speed of SDRAMs continues to limit the speed at which processors may complete operations requiring access to system memory.

A solution to this operating speed disparity has been proposed in the
25 form of a packetized memory device known as a SLDRAM memory device. In the SLDRAM architecture, the system memory may be coupled to the processor, either directly through the processor bus or through a memory controller. Rather than requiring that separate address and control signals be provided to the system memory, SLDRAM memory devices receive command packets that include both control and
30 address information. The SLDRAM memory device then outputs or receives data on a

data bus that may be coupled directly to the data bus portion of the processor bus. A master clock signal transmitted to each memory device is used to synchronize data transfer between the processor and memory device and also serves as a basis from which to generate internal clock signals coordinating internal memory operations.

5 One of the factors limiting the access speed of SDRAM memory devices is the speed at which the command buffer of each device can store and process the command packets. The processing speed of the command buffer is dependent on the control of the relative timing between transmission of the command packets from the processor and an internal clock signal ICLK of the memory device used to trigger a
10 latch in the command buffer to capture the command signals. Both the command signals and the ICLK signal are delayed relative to receipt of the command packet on a command bus and a command clock signal CMDCLK. Furthermore, the amount of the delay is highly variable, and it is difficult to control. If the delay of the internal clock signal ICLK cannot be precisely controlled, it may cause the latch in the command
15 buffer to latch invalid command signals. Thus, the speed at which command packets can be applied to the memory device is limited by the delays in the memory device. Similar problems exist for other control signals in the memory device that control the operation of the memory device during each clock cycle, such as latching of data in the memory device and in a memory controller.

20 Consequently, the operation of a SDRAM memory architecture necessitates the generation of a sequence of clock signals having predetermined phases relative to a master clock signal. Phase-locked and delay locked loops have been employed to ensure the precise phase relationship between clock signals. In such a closed loop, there is typically a phase detector receiving two clock signals, and a
25 voltage controlled delay circuit through which one clock signal passes. The voltage controlled delay circuit receives control signals from the phase detector that are used to adjust the variable delay value in order to establish a predetermined phase relationship between the two clock signals. For example, where the desired phase relationship between two clock signals is zero degrees, the phase detector will detect any phase
30 difference between the two clock signals and generate a control signal that is

transmitted to the voltage controlled delay circuit. The delay circuit will adjust the delay value according to the control signal until the clock signal passing through the voltage controlled delay circuit is synchronized with the other clock signal. The clock control circuitry in an SDRAM is described in greater detail in U.S. Patent
 5 Application Nos. 08/879,847, 08/890,055, 08/933,324, 08/994,461, 09/146,716, and 09/150,079, which are incorporated herein by reference.

A single phase detector connected to a CMOS inverter has been used as a means of providing a control signal to the above-described voltage controlled delay circuits. As shown in Figure 1, clock signals CLK1 and CLK2 are applied to two pulse
 10 generating circuits 11, 12, each of which includes a NAND gate 16 receiving a respective clock signal directly and through three series connected inverters 18, 20, 22. The output of each pulse generating circuit 11, 12 set and reset a flip-flop 26 formed by cross-coupled NAND gates 28, 30. A single output of the flip-flop 26 is connected to the gates of an inverter 36 formed by a PMOS transistor 38 and an NMOS transistor 40.
 15 A current source 44 supplies current to the source of the PMOS transistor 38, and a current sink 46 draws current from the source of the NMOS transistor 40. When the output from the flip-flop 26 is low, the PMOS transistor 38 is turned ON and the NMOS transistor 40 is turned OFF. In this condition, a conductive path is created for the current source 44 to couple current to a capacitor 48. A control signal VOUT is
 20 generated by the capacitor 48. When the current source 44 is applying current to the capacitor 48, the voltage of the control signal VOUT increases linearly. In the alternative case where the output from the flip-flop 26 is high, the PMOS transistor 38 is switched OFF and the NMOS transistor 40 is switched ON. The current sink 46 is then coupled to the capacitor 48 to draw current from the capacitor 48. The voltage of
 25 the control signal VOUT then decreases linearly. As a result, the control signal VOUT has a sawtooth waveform component.

The problem with using a single phase detector connected to an inverter 36, as shown in Figure 1, is that even after the voltage controlled delay circuit has been adjusted so the clock signals have the predetermined phase relationship, the circuit will
 30 nevertheless continue to generate a sawtooth ripple voltage at its output. The sawtooth

waveform component of the control signal VOUT is transmitted to the voltage controlled delay circuit (not shown in Figure 1), which is forced to constantly adjust the delay value, and consequently, the phase relationship between the two clock signals CLK1 and CLK2. The closed loop system will oscillate around a center-point and
5 continue to "hunt" for the optimum control voltage value.

The result is a "phase jitter" imparted to clock signals used to latch command and data signals. Although the phase jitter introduced by the sawtooth ripple voltage may be acceptable in some applications, in high speed memory applications where the clock frequencies are high and the need to control the phase relationship
10 between clock signals is critical, the clocks signals may fail to correctly latch command and data signals.

To accommodate the problems associated with the sawtooth ripple, the memory system designer may relax the timing requirements of the memory system by slowing down the clock frequencies and reducing the operating speed of the memory
15 device. However, this approach defeats the primary purpose of developing high speed memory systems. Therefore, there is a need for a phase detector that generates a control signal that does not vary when the input clock signals have been adjusted to a predetermined phase relationship.

SUMMARY OF THE INVENTION

20 A phase detector used for generating a control signal based on the phase relationship between two clock signals. The phase detector includes two phase detector circuits that each provide to a charge pump or a phase dependent signal source select signals based on the phase relationship of the clock signals. The charge pump receives the select signals and produces a current output signal according to combination of the
25 select signals from the phase detector circuits. The current output signal may be converted into a control signal by connecting a capacitor to the output of the charge pump. Significantly, the phase detector produces a non-varying control signal when the two clock signals have a predetermined phase relationship. The use of two phase

detectors and the charge pump to generate control signals avoids the presence of a sawtooth ripple voltage at the output of the charge pump.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a logic diagram of a conventional phase detector circuit and
5 charge pump.

Figure 2 is a block diagram illustrating a phase detector including two phase detector circuits connected to a charge pump.

Figure 3 is a logic diagram of the phase detector circuits of Figure 2 in accordance with an embodiment of the present invention.

10 Figure 4, comprising Figures 4a-4c, is a timing diagram showing several of the waveforms present in the phase detector circuits.

Figure 5 is a logic diagram of the charge pump of Figure 2 in accordance with an embodiment of the present invention.

15 Figure 6 is a block diagram of a clock generator circuit using an embodiment of the phase detector of Figure 2.

Figure 7 is a block diagram of a computer system using a plurality of DRAMs, each of which includes the phase detector of Figure 2.

DETAILED DESCRIPTION OF THE INVENTION

An embodiment of a phase detector 10 in accordance with the present
20 invention is illustrated in Figure 2. The phase detector includes two single-to-dual signal converters 102, 104 that receive input clock signals CLK1, CLK2, respectively, and produce complementary clock signals CLK1*, CLK2* and non-complementary clock signals CLK1, CLK2 based on the original input clock signals. The single-to-dual signal converters may be implemented using a variety of designs known to one
25 skilled in the art. For example, an inverter and transfer gate having the same propagation delay connected in parallel will produce a complementary and a non-complementary signal from an input clock signal. The CLK1, CLK1*, CLK2, CLK2* signals are transmitted to two phase detector circuits 100, 101 that produce select

signals OUT1, OUT1*, OUT2, OUT2* and transmits them to a phase dependent signal source or a charge pump 200 via signal lines 106, 107, 108, 109, respectively. The charge pump 200 in turn generates an output current IOUT according to the OUT1, OUT1* signals from the phase detector circuit 100, and the OUT2, OUT2* signals from the phase detector circuit 101. The IOUT current may be converted into a control signal V_0 by capacitor 20 connected to an output 280, and the control signal V_0 may be used to adjust the delay value of a voltage controlled delay circuit.

Figure 3 illustrates an embodiment of the phase detector circuits 100, 101 in greater detail. Each detector circuit 100, 101 consists of two signal transition detectors 110, 120 that generate a trigger pulse upon detecting a low-to-high transition of an input clock signal it receives. Each detector circuit 100, 101 also includes a dual output flip-flop 150 that is set and reset by the trigger pulses it receives from the signal transitions detectors 110, 120. Each flip-flop 150 is formed by a pair of cross-coupled NAND gates 152, 154. The output from the phase detector circuits 100, 101 are applied to NAND gates 165, 166 through inverters 161-164. Each NAND gate 165, 166 also receives a respective input signal from the flip-flop 150. The combination of inverters 161-164 and NAND gates 165, 166 creates a buffer circuit that will cause the respective NAND gates 165, 166 to immediately switch upon receiving a trigger pulse from the respective signal transition detector, as explained below.

Each phase detector circuit 100, 101 receives a pair of active-low control signals, SETA*, RSTA*, and SETB*, RSTB*, respectively. The SETA* and RSTA* signals are applied to the phase detector circuit 100, and the SETB* and RSTB* signals are applied to the phase detector circuit 101. The control signals are generated by a control circuit (not shown), and are used to put each phase detector circuit 100, 101 into a predetermined state. Signal RSTA* is provided directly to the NAND gate 152 of the flip-flop 150 and the NAND gate 113 of the signal transition detector 120, and signal SETA* is provided directly to the NAND gate 154 of the flip-flop 150 and the NAND gate 113 of the signal transition detector 110. The RSTB*, SETB* signals are similarly provided to the respective NAND gates of the phase detector circuit 101.

To illustrate the operation of the control signals, consider the effect the SETA*, RSTA* signals have on phase detector circuit 100. During normal operation of the phase detector circuit 100, the RSTA* and SETA* signals are both high. In this situation, the NAND gates 113 of the signal transition detectors 110, 120 behave as
 5 inverters, and the NAND gates 152, 154 behave as simple two input NAND gates. However, when the SETA* signal goes low, the output of the NAND gates 113a and 154 are forced high. Consequently, the NAND gate 118a outputs a high to the NAND gate 152, and the NAND gate 118b outputs a low that sets the flip-flop 150. The OUT1 signal is then forced high and the OUT1* signal is forced low. In a similar manner,
 10 when the RSTA* signal goes low, the flip-flop 150 is reset so that the OUT1 signal is forced low and the OUT1* signal is forced high. Control signals RSTB* and SETB* operate in the same manner for the phase detector circuit 101 by forcing the OUT2, OUT2* signals to a predetermined state when active. To simplify the explanation of the operation of the signal transition detectors 110, 120, it will be assumed that the SETA*,
 15 RSTA*, SETB*, and RSTB* signals are inactive (*i.e.*, at a high signal level).

In operation, the CLK1, CLK2 signals are initially low thereby applying a low signal directly to one input of the NAND gates 118b, 118c and causing the inverters 116a, 116d to apply low signals to the other input of the NAND gates 118a, 118d, respectively. Thus, the NAND gates 118a-d initially output a high signal. When
 20 the respective clock signal goes high, *e.g.*, the CLK1 signal, the NAND gate 118c outputs a low signal until the high signal has propagated through inverter 112c, NAND gate 113c, and through series inverters 114c, 115c, 116c. The inverter of 116c then applies a low signal to the NAND gate 118c, thereby causing the output of the NAND gate 118c to again go high. Thus, the signal transition detector 110 outputs a low-going
 25 pulse responsive to the CLK1 signal. The low-going pulse has a width equal to the total propagation delay through inverter 112, NAND gate 113, and series inverters 114, 115, 116. The signal transition detectors 110, 120 in the phase detector circuit 100, and the signal transition detector 120 in the phase detector circuit 101, each operate in the same manner to output a low pulse responsive to the rising edge of the clock signal to which
 30 it is connected.

The low-going pulse from each of the signal transition detectors 110, 120 sets or resets the flip-flops 150. More specifically, each flip-flop 150 is set by each pulse from the respective signal transition detector 110, thereby causing the NAND gate 152 to output a high signal and the NAND gate 154 to output a low signal. Each flip-flop 150 is reset by each pulse from the respective signal transition detector 120, thereby causing the NAND gate 152 to output a low signal and NAND gate 154 to output a high signal. The output of NAND gates 152, 154 are then inverted by NAND gates 165, 166, respectively, of the buffer circuit to provide the OUT1, OUT1*, OUT2, OUT2* signals to the charge pump 200. As a result, the OUT1 signal is high during the period between the rising edge of the CLK2 signal and the falling edge of the CLK1 signal (*i.e.*, the rising edge of the CLK1* signal). In a similar manner, the OUT2 signal generated by the detector circuit 101 is high during the period between the falling edge of the CLK2 signal (*i.e.*, the rising edge of the CLK2* signal) and the rising edge of the CLK1 signal.

To illustrate the operation of the phase detector circuits 100, 101, consider three situations: first, where CLK1 and CLK2 are in phase; second, where CLK1 is leading CLK2 by ϕ ; and third, where CLK1 is lagging CLK2 by ϕ .

The phase relationship when the CLK1 and CLK2 signals are in phase is illustrated in Figure 4a. As explained above, the OUT1 signal from the phase detector circuit 100 switches from low to high on the rising edge of the CLK2 signal, and from high to low on the falling edge of the CLK1 signal. Also, the OUT2 signal from the phase detector circuit 101 switches from low to high on the falling edge of the CLK2 signal, and from high to low on the rising edge of the CLK1 signal. Since the CLK1 signal is shown in Figure 4a as being in phase with the CLK2 signal, the duty cycles of the OUT1 and OUT2 signals are both 50 percent, and the two signals will never be at the same logic level simultaneously.

Now consider the case where CLK1 is leading CLK2 by ϕ , as shown in Figure 4b. When CLK1 is leading CLK2 by ϕ , the OUT1 signal from the phase detector circuit 100 and the OUT2 signal from phase the detector circuit 101 have duty cycles less than 50 percent, and may be at a low logic level simultaneously. Finally,

consider the case where CLK1 is lagging CLK2 by ϕ , as shown in Figure 4c. When CLK1 is lagging CLK2 by ϕ , the resulting OUT1 and OUT2 signals from phase detector circuits 100 and 101, respectively, have duty cycles greater than 50 percent. Thus, the OUT1 and OUT2 signals may be at a high logic level simultaneously.

5 The OUT1, OUT1*, OUT2, OUT2* signals are transmitted from the phase detector circuits 100, 101 on signal lines 106, 107, 108, 109, respectively, to the input of a charge pump, such as a charge pump 200, as shown in Figure 5. The charge pump 200 includes a charging circuit 205, a current source 270, and a current sink 272. The function of the charging circuit 205 is to direct the current of the current source 270
10 and current sink 272 into, or out of the capacitor 20 (Figure 2), respectively, depending upon the relative duty cycles of the OUT1 and OUT2 signals. Significantly, no current flows into or out of the capacitor 20 when the CLK1 and CLK2 signals are in phase. Thus, the resulting control signal may have virtually no ripple when the phase detector
15 10 (Figure 2) is used in a voltage controlled delay circuit, as explained above. A clock signal generated using the voltage controlled delay circuit has significantly less phase jitter compared to a clock signal generated by a voltage controlled delay circuit using the phase detector of Figure 1.

 The charge pump 200 includes transistors 245-248 on the left leg of the charging circuit 205 to form a compensation circuit 206 to compensate for current and
20 voltage changes in a current driving circuit formed by transistors 243, 244, 249, and 250 on the right leg of the charging circuit 205. The compensation circuit 206 is provided so that the voltage across the charging circuit 205 is relatively constant during operation, regardless of where the currents of the current source 270 and the current sink 272 are being directed.

25 A voltage follower 260 is connected between the output 280 of the charging circuit 205 and node 262 of the compensation circuit 206. The voltage follower 260 provides a current path from the current source 270 to ground when current from the current sink 272 is being directed out of the capacitor 20. The voltage follower 260 also provides a current path from the current sink 272 to the positive
30 supply when current from the current source 270 is directed to the output 280. As will

be explained in greater detail below, both of these situations occur where the CLK1 and CLK2 signals are not in phase. As a result, the current through the charging circuit 205 is through two PMOS transistors and two NMOS transistors when current is being directed into or out of the capacitor 20, namely, the transistors 242, 244, 247, 251 or the transistors 241, 245, 250, 252. Similarly, in the situations where no current is being directed into or out of the capacitor 20, the current through the charging circuit 205 is also through two PMOS transistors and two NMOS transistors, namely, the transistors 242, 243, 248, 251 or the transistors 241, 246, 249, 252. Consequently, the operating points of the active transistors will remain relatively constant, and any capacitive charge pumping on the internal nodes of the charging circuit 205 will be minimized. It will be appreciated by one ordinarily skilled in the art that the transistors of the charging circuit 205 must be scaled accordingly.

To illustrate the operation of the charge pump 200 in conjunction with the phase detector circuits 100, 101, consider again the three situations that were described earlier: where CLK1 and CLK2 are in phase; where CLK1 is leading CLK2 by ϕ ; and where CLK1 is lagging CLK2 by ϕ .

As shown in Figure 4a, when OUT1 is low and OUT2 is high the current provided by the current source 270 and sunk by the current sink 272 bypasses the output 280 of the charging circuit 205 and simply flows through the transistors 242, 243, 248, 251 (indicated in Figure 4a as "0(A)"). Similarly, when OUT1 is high and OUT2 is low current flows from the current source 270 to the current sink 272 through the transistors 241, 246, 249, 252 (indicated in Figure 4a as "0(B)"). In either case, the charging circuit 205 does not charge or discharge the capacitor 20 so the voltage on the capacitor 20 remains constant.

As shown in Figure 4b, when both the OUT1 and OUT2 signals are low, the charging circuit 205 directs the current provided by the current source 270 through the PMOS transistors 242, 244 to charge the capacitor 20 (indicated in Figure 4b as "I+"). A path for the current from the current sink 272 is provided through the voltage follower 260 and the NMOS transistors 247, 251. During the time the OUT1 and OUT2 signals are at different logic levels, the charging circuit 205 does not charge or

discharge the capacitor 20 so the voltage on the capacitor 20 remains constant, as was previously explained (indicated in Figure 4b as 0(A) or 0(B)).

As shown in Figure 4c, when both the OUT1 and OUT2 signals are high, the NMOS transistors 250, 252 provide a conductive path for the current sink 272 to
 5 sink current from the capacitor 20 (indicated in Figure 4c as "I-"). A current path for the current of the current source 270 is provided through the voltage follower 260 and the PMOS transistors 241, 245. As mentioned before, whenever the OUT1 and OUT2 signals are at different logic levels, the charging circuit 205 does not charge or discharge the capacitor 20 so the voltage on the capacitor 20 remains constant (indicated
 10 in Figure 4c as 0(A) or 0(B)).

Any change in the control voltage V_0 depends upon whether current is flowing into or out of the capacitor 20, as explained above. When the CLK1 and CLK2 signals are in phase, as illustrated in Figure 4a, OUT1 and OUT2 never have the same logic level so no current flows either into or out of the capacitor 20. In contrast, when
 15 the CLK1 signal and the CLK2 signal have different phases, OUT1 and OUT2 are both high or both low for a portion of each cycle. As illustrated in Figure 4b, when the CLK1 signal leads the CLK2 signal, OUT1 and OUT2 are low for more than 50 percent of each cycle so that OUT1 and OUT2 are both low for a portion of each cycle. As a result, as explained above, current flows into the capacitor 20, thereby increasing the
 20 control voltage V_0 . Similarly, when the CLK1 signal lags the CLK2 signal as illustrated in Figure 4c, OUT1 and OUT2 are high for more than 50 percent of each cycle so that OUT1 and OUT2 are both high for a portion of each cycle. As a result, current flows out of the capacitor 20, thereby decreasing the control voltage V_0 .

The discussion of the phase detector 10 has so far only considered the
 25 case where the CLK1 and CLK2 signals are adjusted so that they are approximately in phase. However, the phase detector 10 may be modified to produce a control signal that adjust the CLK1 and CLK2 signals to have a 180 degrees phase relationship. As shown in Figure 3, the CLK1* and CLK2 signals are transmitted to nodes 95, 96 of the phase detector circuit 100, while the CLK1 and CLK2* signals are transmitted to nodes 97, 98
 30 of the phase detector circuit 101, resulting in a phase detector that generates a non-

varying control signal when the CLK1 and CLK2 signals are in phase. However, when the CLK1 and CLK1* signals are reconnected to the nodes 95 and 97, respectively, or the CLK2 and CLK2* signals are reconnected to the nodes 98 and 96, respectively, the phase detector circuits 100, 101 transmit the OUT1, OUT1*, OUT2, OUT2* signals to
 5 the charge pump 200 so that the phase detector 10 generates a non-varying control signal when the CLK1 and CLK2 signals have a 180 degree phase relationship.

The current source 270 and the current sink 272 of the charge pump 200 may be of any current source circuit known in the art. In a preferred embodiment, a high-swing cascode current mirror, as described in "CMOS Circuit Design, Layout, and
 10 Simulation," published by IEEE Press, is used for both the current source 270 and the current sink 272. The use of this particular current source is meant for illustrative purposes only, and is not intended to limit the scope of the present invention.

The charge pump 200 (Figure 2) has been described with respect to the embodiment illustrated in Figure 5. However, an integrator circuit using an operational
 15 amplifier may also be used for the charge pump 200. Such an integrator circuit is formed by coupling a capacitor across the output of the operational amplifier and the inverting input. The OUT1 and OUT2 signals generated by the phase detector circuits 100 and 101 are applied through two resistors of equal resistance to the non-inverting input of the operational amplifier, and the OUT1* and OUT2* signals are applied
 20 through two resistors of equal resistance to the inverting input. The resulting charge pump will generate increasing and decreasing control signals when the CLK1 and CLK2 signals are not in phase, and generate a control signal with virtually no ripple when the clock signals are in phase.

Shown in Figure 6 is a block diagram of a clock generator circuit that
 25 may be used in packetized DRAMs to provide a sequence of clock signals that have predetermined phases relative to a master clock signal. The clock generator circuit contains a first delay-locked loop 301 and a second delay-locked loop 302, each having a phase detector 10 of Figure 2. A multiplexer 330 having a plurality of output lines coupled to respective clock drivers 314a-n may be coupled to the first delay-lock loop
 30 301 to couple one of the clock signals produced by the multi-tap voltage controlled

delay circuit 310 to a clock output terminal 316 for use, for example, with a latch 340 to latch command data CMD DATA in packetized DRAM. The multiplexer 330 couples the input of each of the clock drivers 314a-n to any one of the clock signals produced by the multi-tap voltage controlled delay circuit 310.

5 The first delay-locked loop includes a multi-tap voltage controlled delay circuit 310 and a first phase detector 10a. The multi-tap voltage controlled delay circuit 310 generates a sequence of clock signals on output lines 312a-312n that are increasingly delayed from a first clock signal on line 312a to a last clock signal on line 312n. Two of the clock signals, preferably the first and last clock signals, are locked to
10 each other using the delay-locked loop 301 so that they have a predetermined phase with respect to each other. For example, the first clock signal on line 312a and the last clock signal on line 312n may be locked so that they are the inverse of each other, that is, the predetermined phase relationship is 180 degrees from each other. Alternatively, the predetermined phase relationship could be 360 degrees so the first and last clock
15 signals are in phase. The first phase detector 10a compares the phase of the clock signals on lines 312a and 312n and generates the first control signal as a function of the phase difference therebetween. The first control signal is provided to the multi-tap voltage-controlled delay circuit 310 on line 311 to adjust the relative delay between the clock signal on line 312a and line 312n. The phase detector 10a will continue to
20 provide the first control signal until the first and last clock signals have obtained the predetermined phase relationship.

Likewise, the second delay-locked loop 302 includes a second voltage controlled delay circuit 320 and a second phase detector 10b. A second delay-locked loop locks a clock signal from the multi-tap voltage controlled circuit 310 to a master
25 clock signal CMD CLK on line 305 so that the increasingly delayed clock signals of the multi-tap voltage controlled delay circuit 310 have phase delays with respect to the CMD CLK signal. The clock signal from the multi-tap voltage controlled circuit 310 is provided to an input of the second phase detector 10b through a simulated multiplexer 317 and a clock driver 318. The relative phase delays of the simulated multiplexer 317
30 and the clock driver 318 are nearly identical to that of the multiplexer 330 and the clock

drivers 314a-n. Consequently, the phase detector 10b will receive a clock signal having the same relative phase delay as a clock signal output by the clock drivers 314a-n.

For example, the second delay-lock loop 302 may delay lock the first clock signal on line 312a to the CMD CLK signal so that they have substantially the same phase, that is, the predetermined phase relationship is zero degrees from each other. The voltage controlled delay circuit 320 receives the CMD CLK signal and generates a reference clock signal on line 322 having a delay relative to the CMD CLK signal that is a function of a second control signal on line 321. The clock signal on line 322 is provided to the multi-tap voltage controlled circuit 310 and used to generate the sequence of increasingly delayed clock signals 312a-312n.

The second phase detector 10b compares the phase of the CMD CLK signal to the phase of the first clock signal on line 312a and generates a second control signal as a function of the difference therebetween. The clock signal provided to the phase detector 10b is delayed through the simulated multiplexer 317 and the clock driver 318 approximately the same amount as the clock signals output by the clock drivers 314a-n. The second control signal is used to adjust the delay value of the voltage controlled delay circuit 320. The second control signal is provided by the second phase detector 10b until the CMD CLK signal and the first clock signal from the multi-tap voltage controlled circuit 310 have obtained the predetermined phase relationship. The clock generator circuit of Figure 6 is described in greater detail in U.S. Patent Application No. 08/879,847, which, as mentioned above, has been incorporated herein by reference.

A computer system using the phase detector 10 of Figure 2 in each of a plurality of packetized DRAMs 401 is shown in Figure 7. With reference to Figure 7 the computer system 400 includes a processor 402 having a processor bus 404 coupled to three packetized dynamic random access memory or SDRAMs 401a-c. The computer system 400 also includes one or more input devices 410, such as a keypad or a mouse, coupled to the processor 402 through a bus bridge 412 and an expansion bus 414, such as an industry standard architecture ("ISA") bus or a Peripheral component interconnect ("PCI") bus. The input devices 410 allow an operator or an electronic

device to input data to the computer system 400. One or more output devices 420 are coupled to the processor 402 to display or otherwise output data generated by the processor 402. The output devices 420 are coupled to the processor 402 through the expansion bus 414, bus bridge 412 and processor bus 404. Examples of output devices 5 420 include printers and video display units. One or more data storage devices 422 are coupled to the processor 402 through the processor bus 404, bus bridge 412, and expansion bus 414 to store data in or retrieve data from storage media (not shown). Examples of storage devices 422 and storage media include fixed disk drives floppy disk drives, tape cassettes and compact-disk read-only memory drives.

10 In operation, the processor 402 communicates with the memory devices 401a-c via the processor bus 404 by sending the memory devices 401a-c command packets that contain both control and address information. Data is coupled between the processor 402 and the memory devices 401a-c, through a data bus portion of the processor bus 404. Although all the memory devices 401a-c are coupled to the same 15 conductors of the processor bus 404, only one memory device 401a-c at a time reads or writes data, thus avoiding bus contention on the processor bus 404. Bus contention is avoided by each of the memory devices 401a-c and the bus bridge 412 having a unique identifier, and the command packet contains an identifying code that selects only one of these components.

20 The computer system 400 also includes a number of other components and signal lines which have been omitted from Figure 7 in the interests of brevity. For example, as explained above, the memory devices 401a-c also receive a command or master clock signal to provide internal timing signals, a data clock signal clocking data into and out of the memory device 401a-c, and a FLAG signal signifying the start of a 25 command packet.

From the foregoing it will be appreciated that, although specific embodiments of the invention have been described herein for purposes of illustration, various modifications may be made without deviating from the spirit and scope of the invention. For example, the signal transition detector circuits 110, 120 have been 30 described as generating a negative trigger pulse upon detecting a transition of an input

signal. However, using cross-coupled NOR gates instead of cross-coupled NAND gates for the flip-flop 150 allows the use of a signal transition detector that generates a positive trigger pulse. Also, the charge pump 200 shown in Figure 5 is described as generating a current output signal IOUT having a positive polarity when CLK1 is leading CLK2, and having a negative polarity when CLK1 is lagging CLK2. However, the connection of the OUT1, OUT1*, OUT2, OUT2* signals to the transistors 241-252 of the charging circuit 205 may be modified so that the IOUT signal will have a negative polarity when CLK1 is leading CLK2, and have a positive polarity when CLK1 is lagging CLK2. Accordingly, the invention is not limited except as by the appended claims.